



SOLASAI

**Beyond Measurement: Lessons in Mitigating Algorithmic Bias
in Consumer Finance and Healthcare**

Nicholas Schmidt

Workshop on Mitigating AI Bias in Context, August 31, 2022

Introduction

- Qualifications and Experience
- Mitigation Works
- The Measurement Problem
- Existing Frameworks for Disparity Mitigation
- Models versus Model Systems

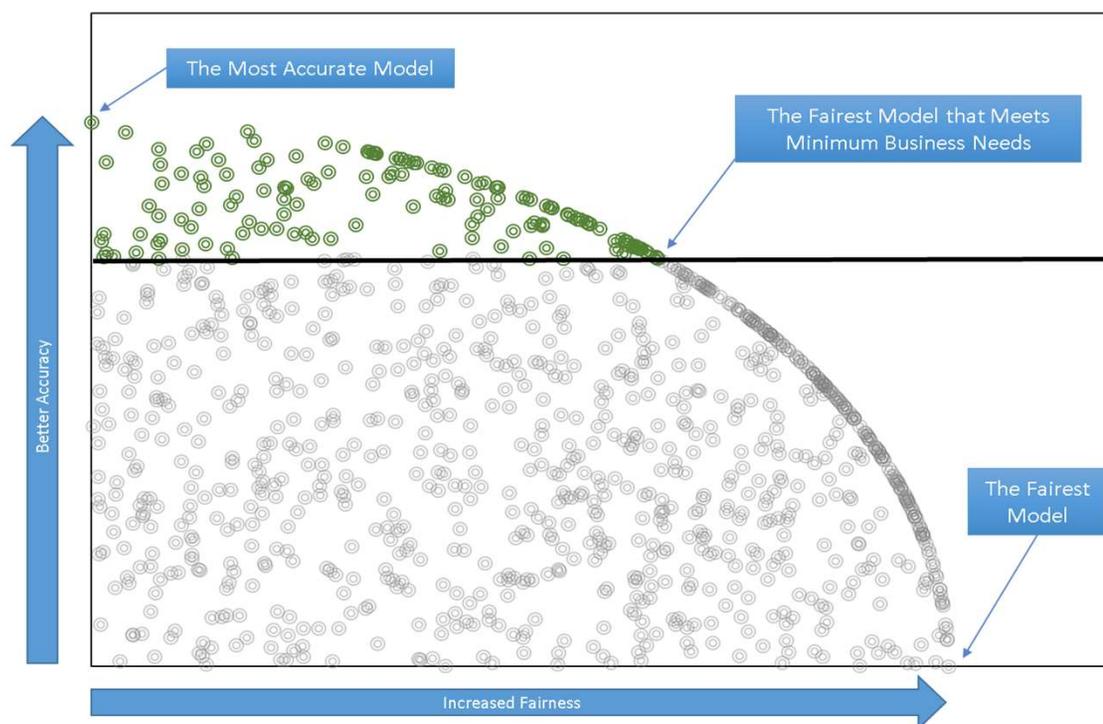
Background

- **Nicholas Schmidt**
 - 20+ years of experience applying concepts from statistics and economics to questions of law and regulatory compliance.
- **CEO, SolasAI**
 - SolasAI software *measures* and *mitigates* discrimination risk.
 - Prominent U.S. lenders, insurers, and health insurance companies are using SolasAI to assess and mitigate discrimination risk.
- **AI Practice Leader, BLDS, LLC**
 - We are the fair lending analytics advisors to lenders that represent over 70% of credit cards issued in the United States.
 - We are regularly engaged by regulators and courts to provide guidance on discrimination risk in algorithms.

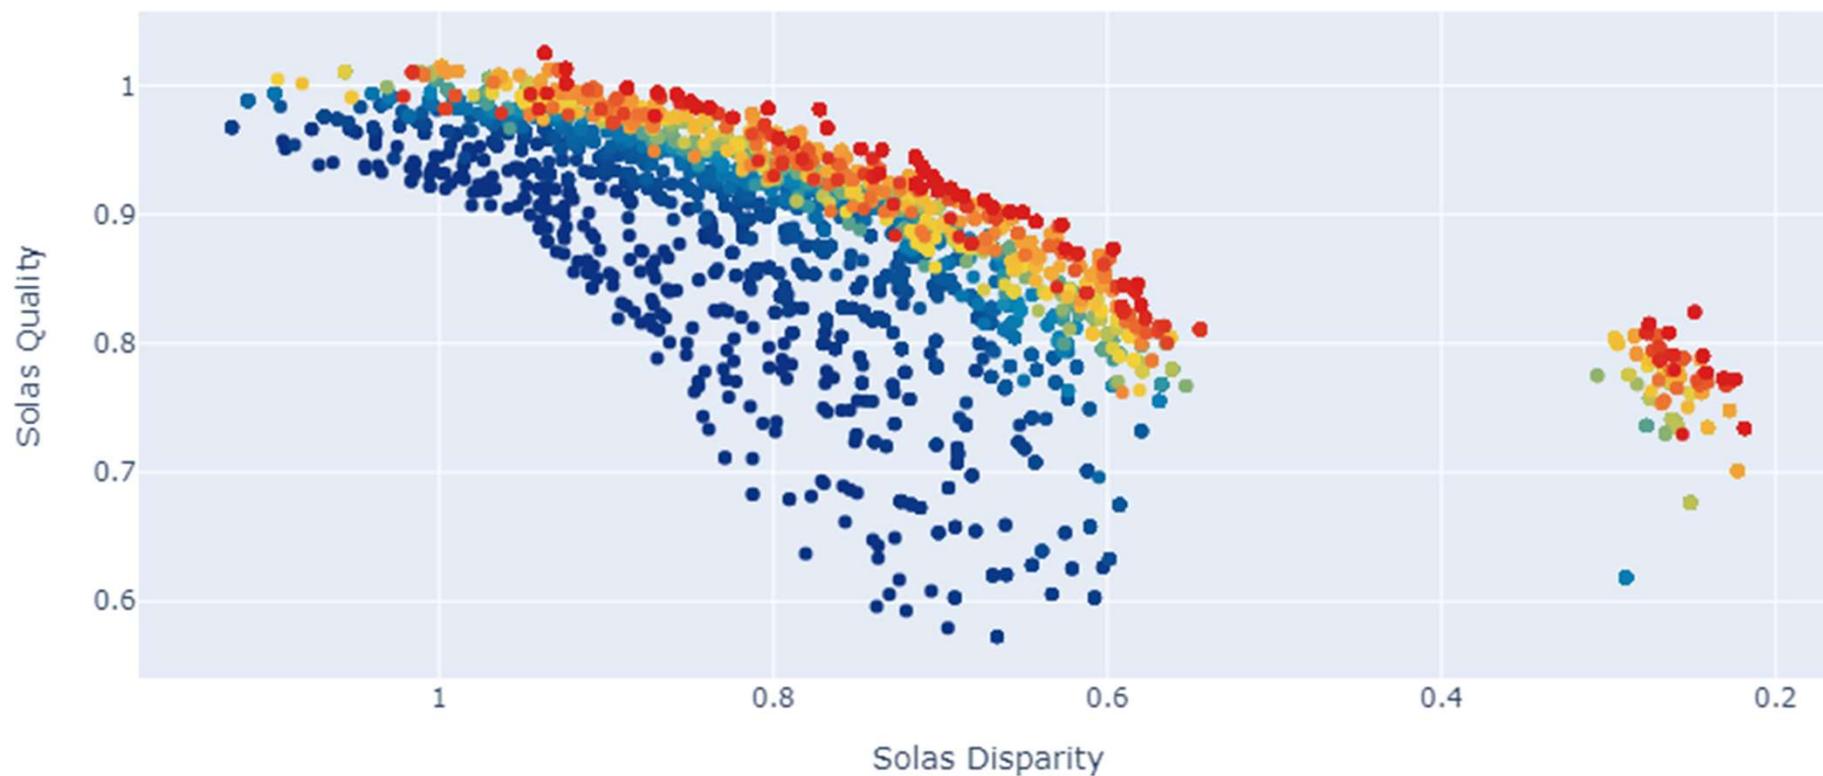


Fixing Algorithmic Discrimination: The Pareto Frontier

- Most organizations want to minimize algorithmic bias.
- But they exist in a competitive environment where small drops in model accuracy may lead to large relative losses.
- The solution is to do the best the organization can do.



Mitigating Discrimination in Machine Learning: A Real-World Example



The Measurement Problem as Pretext

- Fairness metrics can be contradictory:
 - *"If we can't define fairness, how can we do anything about it?"*
 - Does this mean we should just let discrimination continue?
 - Of course not.
- What is missing:
 - Use intelligence, common sense, and existing public policies to guide decisions.
 - Recognize that there is rarely one "correct" model.
 - Understand that doing better is better than doing nothing at all.

An Existing Regulatory Framework for Fairness Mitigation

- First, Remove **disparate treatment**.
- Second, Follow the framework outlined in the “**burden-shifting test:**”
 1. Does the model lead to a **disparate impact**? (are unconditioned outcomes different?)
 2. Does the model have a **valid business justification**? Is it “empirically derived, demonstrably and statistically sound?”
 3. **Are there alternative models that are fairer, but maintain predictive ability?**
- The multiplicity of good models means we almost always can find alternative specifications that reduce discrimination while keeping a “good” model.

Addressing Discrimination in Models

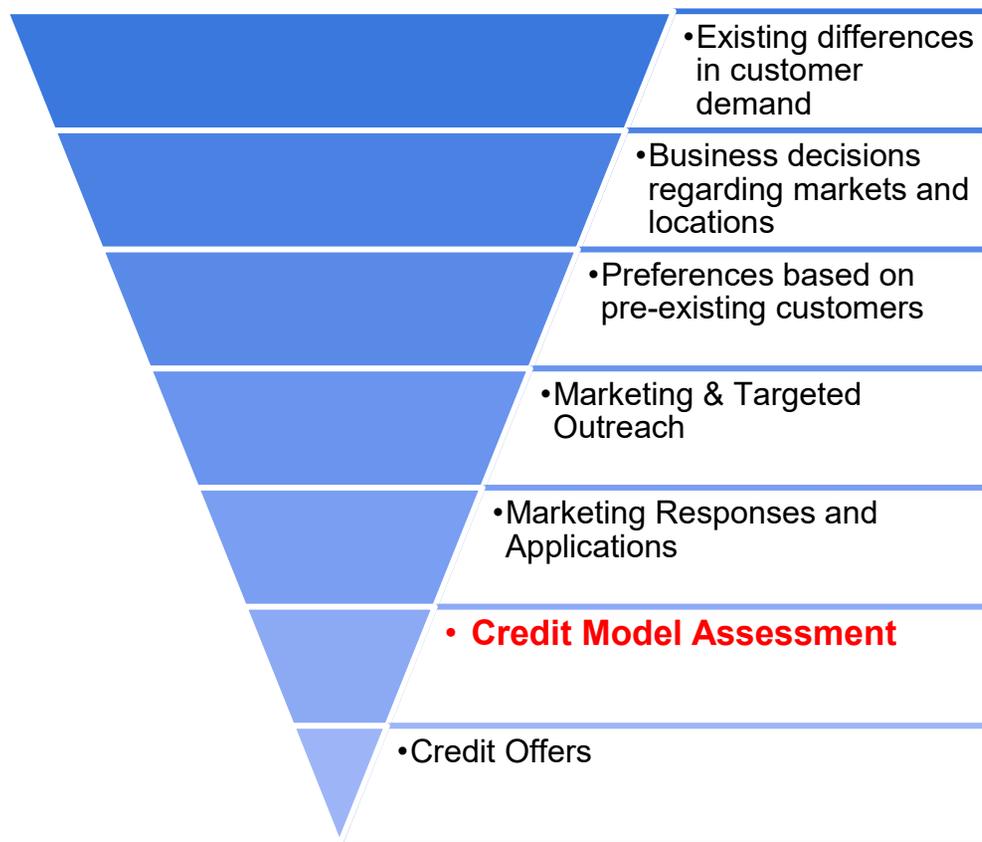
Addressing and mitigating discrimination risk requires effort throughout the model pipeline, from data ingestion to model monitoring.

These steps include:

1. Limit data availability
2. Review data and model usage for legal, ethical, and reputational risk
3. Unless justified and legal, do not use protected class data
4. Test the model for disparate impact and, if necessary, proxy feature(s)
5. Search for less discriminatory alternative models
6. Monitor for drift - especially usage drift

The Problem of Model Systems

- Addressing algorithmic discrimination is essential.
- But algorithms may play a small role in a model system.
- Model systems can and should be evaluated for the potential to mitigate discrimination.



Pitfalls in Fairness Analyses

- **Putting garbage or risky data in a model**
 - Are the features discriminatory?
 - How are you using non-standard data (e.g., second look models)?
- **Not measuring disparities considering actual outcomes**
 - Unless you give loans to people with a 50% probability of default, then you should not use a 50% cutoff to measure disparities.
- **Getting compliance advice from people with no compliance or business experience**
 - Great work is being done on fairness in AI. But do not ignore real-world aspects, like regulatory compliance.
- **Running the cool new de-biasing algorithm**
 - What are the methods truly accomplishing?
 - Are they compliant and legal?
- **Getting too attached to the model you've chosen**
 - There are almost certainly many other models that are similar
 - Assuming causality
 - Assuming a 0.001 increase in AUC is meaningful



Thank You

Nicholas Schmidt

nicholas.schmidt@solas.ai

<https://www.linkedin.com/in/nickpschmidt>